

PATENT
Attorney Docket No. 10354*1

UNITED STATES PATENT APPLICATION

of

John R. HUMPHREY, James P. DURBANO, and Dennis W. PRATHER

for

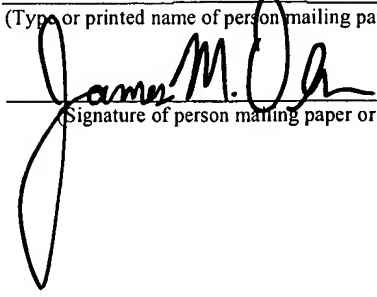
**HARDWARE IMPLEMENTATION OF THE
PSEUDO-SPECTRAL TIME-DOMAIN METHOD**

CERTIFICATE OF EXPRESS MAIL
(37 C.F.R. § 1.10)

I HEREBY CERTIFY THAT THIS PAPER OR FEE IS BEING DEPOSITED WITH THE UNITED STATES POSTAL SERVICE, EXPRESS MAIL POST OFFICE TO ADDRESSEE UNDER 37 C.F.R. § 1.10, BEARING EXPRESS MAIL LABEL NO. **ER 644522188** US ON THIS **8TH** DAY OF **DECEMBER, 2004** AND IS ADDRESSED TO: MAIL STOP PCT, COMMISSIONER FOR PATENTS, P.O. BOX 1450, ALEXANDRIA, VIRGINIA 22313-1450.

James M. Olsen

(Type or printed name of person mailing paper or fee)


(Signature of person mailing paper or fee)

1/987

10/517224
DT05 Rec'd PCT/PTO 08 DEC 2004

CLAIM FOR PRIORITY

The present application is a U.S. National Stage application filed under 35 U.S.C. § 371, claiming priority of International Application No. PCT/US2003/020259, filed June 24, 2003, and U.S. Provisional Patent Application Serial No. 60/390,993, filed June 24, 2002, respectively, under 35 U.S.C. §§ 119 and 365, the disclosures of the above-referenced applications being incorporated by reference herein in their entireties.

BACKGROUND OF THE INVENTION

A. Field of the Invention

The present invention relates generally to hardware accelerators, and, more particularly to a hardware implementation of the pseudo-spectral time-domain (PSTD) method.

B. Description of the Related Art

Since the advent of the modern computer, a great deal of effort has gone into the development of numerical algorithms for the rigorous solution of electromagnetic problems. Today, popular numerical approaches include the finite-element method (FEM), method of moments (MOM), modal expansion techniques, boundary integral methods, and time-domain methods such as the finite-volume time-domain (FVTD), multi-resolution time-domain (MRTD), and the finite-difference time-domain (FDTD) methods. Each of these algorithms possesses clear advantages and disadvantages depending on the specific application. One algorithm in particular, the pseudo-spectral time-domain (PSTD) method shows particular promise. In comparison to the above methods, the PSTD technique can require far less memory while maintaining, and in fact improving, the accuracy and versatility of electromagnetic analysis. To this end, recent numerical experiments have confirmed that, for a fixed amount of computational resources, the PSTD method can analyze

problems two to three orders of magnitude larger than an FDTD method with the same level of accuracy.

However, one of the difficulties of the PSTD method is the large number of forward and inverse Fast Fourier Transforms (FFTs and IFFTs) that need to be computed, which significantly
5 slow down the analysis. In comparison to the FDTD method, which has order N computational dependence, the PSTD method has order $M \log N$. Thus, from a pure software point of view, the PSTD method is far less appealing from a computational resources point of view.

Over the last several decades, significant effort has been put into realizing application-specific integrated circuits (ASICs) for application to digital signal processing (DSP). As a result,
10 ASICs are currently available that perform the FFT and IFFT operations in fractions of a microsecond. Thus, the PSTD method is far more attractive to implement in hardware than the FDTD method, due to the wealth of technology that it can leverage. While the FDTD method can also be realized in hardware, it suffers from the fact that it requires extraordinary amounts of computer memory. In comparison, the PSTD method can analyze problems two to three orders of
15 magnitude larger than other full-wave techniques, the FDTD method in particular, with the same level of accuracy.

Despite the advantages of the FDTD acceleration hardware over software-based implementations, the FDTD method requires a tremendous amount of memory. This can greatly limit the size of the problems that are capable of being solved. Because the PSTD method requires
20 fewer samples than the FDTD method (on the order of 1000 times fewer), much larger problems can be solved with the PSTD method given the same amount of memory.

Although PSTD methods are accurate and well defined, current computer system technology limits the speed at which we can perform these operations. To use the PSTD method to solve a non-

trivial problem can take hours, days, weeks, months, etc. Some problems are even too large to be effectively solved due to time constraints.

Thus, there is a need in the art to overcome the limitations of the related art, and to provide for a practical, hardware implementation of the PSTD method.

5

SUMMARY OF THE INVENTION

The present invention solves the problems of the related art by providing a hardware-based PSTD processor that capitalizes on the large advancements in the area of application DSP ASICs. By combining the PSTD algorithm with modern DSP chips and large scale FPGAs, the PSTD
10 processor will be capable of solving very large radiation problems in computation times short enough to enable iterative design. Given this potential, a hardware implementation of the PSTD method offers the ability to design an extraordinary array of electromagnetic problems that heretofore have been impossible.

Further scope of applicability of the present invention will become apparent from the detailed
15 description given hereinafter. However, it should be understood that the detailed description and specific examples, while indicating preferred embodiments of the invention, are given by way of illustration only, since various changes and modifications within the spirit and scope of the invention will become apparent to those skilled in the art from this detailed description. It is to be understood that both the foregoing general description and the following detailed description are exemplary and
20 explanatory only and are not restrictive of the invention, as claimed.

BRIEF DESCRIPTION OF THE DRAWING

The present invention will become more fully understood from the detailed description given

hereinbelow and the accompanying drawing which are given by way of illustration only, and thus are not limitative of the present invention, and wherein:

Fig. 1 is a schematic diagram showing a hardware implementation of the PSTD method in accordance with an embodiment of the present invention.

5

DETAILED DESCRIPTION OF EMBODIMENTS OF THE INVENTION

The following detailed description of the invention refers to the accompanying drawings. The same reference numbers in different drawings identify the same or similar elements. Also, the following detailed description does not limit the invention. Instead, the scope of the invention is defined by the appended claims and equivalents thereof.

10

Equations in the PSTD method take the following form:

$$E_{ab} = AE_{ab} + B \frac{\partial H_c}{\partial b} + CE_{ab}^{inc} \quad (1)$$

where a , b , and c are directions (x, y, or z), A , B , and C are coefficients based on the material properties of the medium, and E_{ab}^{inc} is the incident field associated with the node.

15

Unlike the finite-difference time-domain (FDTD) method, where both spatial and temporal derivatives are represented by finite differences, the PSTD method uses FFTs and IFFTs to calculate the spatial derivatives. For example, recall from basic frequency-domain calculations that the spatial derivative of a field, H_z , in the y -direction can be written as:

$$\frac{\partial H_z}{\partial y} = \text{IFFT}(j * k_y * \text{FFT}(H_z(i, :, k))) \quad (2)$$

20

To compute the spatial derivative of H_z in the y -direction for some node (i, j, k) , all of the values of H_z in the y -direction along the line $(i, ..., k)$ are required. Although this can require a tremendous amount of data to be fetched (depending on the mesh size), this operation only needs to

be performed once for all values that require this derivative along this line. If there are N values in the y -direction, by fetching N values from memory, the spatial derivative in the y -direction for N different mesh points may be computed. Although there is a significant latency associated with the fetch operation, this operation only needs to be performed once per derivative, per timestep.

5 Despite the large amount of data required to perform the FFT/IFFT operations, it is possible to efficiently organize RAM to maximize throughput. For example, if the values of H_z were stored in RAM in increasing- y order, it would be possible to perform a burst-read from RAM. Bursting allows the RAM to fetch contiguous locations very rapidly. Thus, burst reads allow fetching of all of the values necessary to compute the derivative, while, at the same time, maximizing the throughput
10 of the RAM. Because spatial derivatives are required in the x , y , and z -directions, values would need to be stored in the RAM in three different patterns (increasing- x , increasing- y , and increasing- z). This would permit taking advantage of the bursting capabilities when computing any required derivative. Similarly, there will be three different “update” orders. The update order specifies the order in which nodes in the mesh will be updated. Updating nodes in the same order in which they
15 are stored in RAM allows reuse of the spatial derivatives. Once the spatial derivative of a field is known in a specific direction, that value can be used to update any node along that line. Thus, by solving fields along that line, the same value can be used over and over again without incurring the latency of the RAM fetching and the FFT/IFFT operations.

Fig. 1 shows one logical data flow for a PSTD accelerator 10 in accordance with the present
20 invention. Note that there are three parallel datapaths 12, 14, 16. These datapaths 12, 14, 16 are completely independent of one another (except for a memory subsystem 18) and each solves fields in a given direction (x , y , or z). Each of these datapaths 12, 14, 16 is responsible for solving an equation of the form shown in equation 1.

The flow within each computational path begins by determining the spatial derivative needed for the computation. The necessary values must be fetched from the memory subsystem 18 and streamed into an FFT unit 20 that performs the FFT on the values. Depending on the size of the problem being solved and the capabilities of the FFT unit 20, this computation can easily require several thousand cycles. While the FFT is being computed by the FFT unit 20, the other necessary data, including the primary fields, incident fields, and coefficients, can all be fetched from memory subsystem 18 or computed. This entire computational path will be pipelined to hide much of the latency of the FFT operation. As results begin emerging from the FFT unit 20, the results are streamed through a pipelined a complex multiplication unit 22 that solves the multiplication aspects of equation (2) set forth above. At this point, the spatial derivatives will have been computed in the frequency domain. The next step in each datapath 12, 14, 16 is to convert the frequency domain result back into the time domain by means of an IFFT unit 24 that performs an IFFT. IFFT unit 24 will also undergo a several thousand-cycle latency. As results begin emerging from the IFFT unit 24, the results are streamed into a Computation Engine (CE) 26. The CE 26, given the necessary data and the spatial derivative, solves equation (1) set forth above. Once complete, the fields are written back to the memory subsystem 18 by CE 26.

It will be apparent to those skilled in the art that various modifications and variations can be made in the hardware implementation of the PSTD method of the present invention and in construction of the hardware without departing from the scope or spirit of the invention.

For example, although Fig. 1 shows three parallel computational datapaths 12, 14, 16, more or less such datapaths may be provided. Fewer datapaths may be implemented to save on hardware, or extra datapaths may be included to increase parallelism. The actual number of datapaths is mostly dependent upon the capabilities of the memory subsystem 18.

Also, the FFT and IFFT units 20, 24 may or may not be co-located with complex multiplication unit 22 and CE unit 26. FFT and IFFT units 20, 24 may be implemented inside of a field-programmable gate-array (FPGA). However, custom FFT/IFFT chips may be purchased or DSP chips may be programmed to solve the FFT/IFFT equations. The latter configurations require DSP chips external to the FPGA. This would result in a more complex printed circuit board and latency in transferring data into and out of the FPGA. However, DSP chips are extremely fast and would save FPGA resources. However, with very advanced FPGAs already in the market, an implementation where the FFT/IFFT operations are performed inside the FPGA is quite possible.

By implementing the PSTD method in hardware, computational speedup is achieved that allows solving problems much faster than current software-based methods and also allows solving problems that were heretofore unsolvable due to time constraints. The hardware implementation of the PSTD method of the present invention may be used in any field or application utilizing a software method based on a pseudo-spectral time-domain approach.

Other embodiments of the invention will be apparent to those skilled in the art from consideration of the specification and practice of the invention disclosed herein. It is intended that the specification and examples be considered as exemplary only, with a true scope and spirit of the invention being indicated by the following claims.